Volume 06 Issue 02 Published, May 16, 2002 ISSN 1535766X

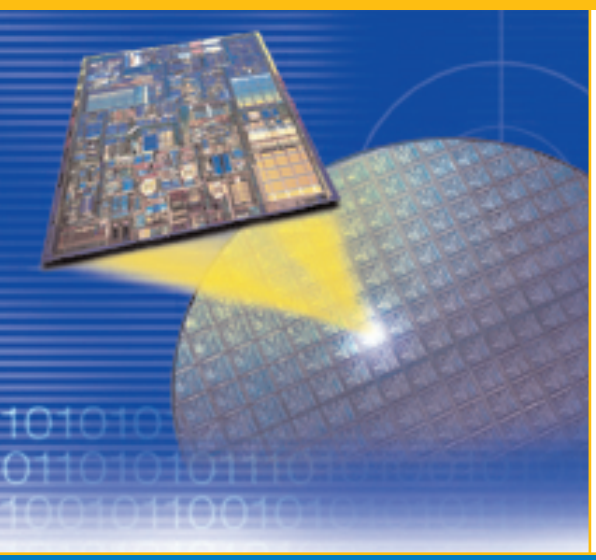

Intel® Technology Journal

Semiconductor Technology and Manufacturing

Process Development and Manufacturing of High-Performance Microprocessors on 300mm Wafers

A compiled version of all papers from this issue of the Intel Technology Journal can be found at:
http://developer.intel.com/technology/itj/index.htm

Process Development and Manufacturing of High-Performance Microprocessors on 300mm Wafers

Sanjay Natarajan, Logic Technology Development, Intel Corporation
Melton Bost, Logic Technology Development, Intel Corporation
Derek Fisher, Logic Technology Development, Intel Corporation
David Krick, Logic Technology Development, Intel Corporation
Chris Kenyon, Logic Technology Development, Intel Corporation
Chris Kardas, Logic Technology Development, Intel Corporation
Chris Parker, Logic Technology Development, Intel Corporation
Robert Gasser, Jr., Logic Technology Development, Intel Corporation

Index words: 300mm, Px60, P1260, 0.13 m, 130nm, Copy Exactly!

ABSTRACT

Over 35 years ago, Moore's Law established the nature of competition in the semiconductor industry by projecting a 2x transistor density improvement approximately every 18 months. Faced with increasingly challenging process technology issues, industry leaders such as Intel have had to achieve increasingly faster yield improvement and volume production ramps to maintain competitiveness. The Copy Exactly! methodology, which has been used since 1992 to transfer technologies and ramp new factories, has been instrumental in allowing Intel to meet these challenges.

The subject of this paper is the successful extension of Copy Exactly! to Intel's first 300mm process technology, P1260, to achieve rapid yield learning and volume production. P1260 replicates Intel's industry-leading 200mm 0.13 m CMOS process in performance, yield, reliability, and density, with SRAM cell sizes below 2 m^2 [1]. Intel has used the Copy Exactly! methodology for several generations with documented success, and we present perhaps the most compelling evidence to date of its utility: accurate replication of an industry-leading 200mm 0.13 m CMOS process on a 300mm wafer size using a completely new process equipment set.

INTRODUCTION

Moore's Law

In 1965, Gordon Moore, then R&D manager at Fairchild Semiconductor and now Chairman Emeritus of Intel Corporation, characterized the rate of progress in the semiconductor industry and arrived at an astounding conclusion: the density of transistors per integrated circuit (IC) had been doubling at regular intervals and would continue to do so indefinitely [2].

The observation, later termed "Moore's Law," has been extremely influential in the semiconductor industry, even to the point of becoming self-fulfilling. Since Moore's Law has accurately predicted past IC growth, it is also viewed as a method for predicting future trends, setting goals for innovation, directing the pace of the technology treadmill, and ultimately defining the nature of industry competition [3].

Delivering the regular progress dictated by Moore's Law in the face of increasingly complex process technologies requires steady improvements in the pace of yield learning and volume manufacturing capability. Figures 1 and 2 illustrate this trend for Intel's process technologies. Figure 1 shows the steadily increasing rate of production ramp for each of the last six process generations. Across these six generations, there has been a 4x increase in the ramp rate, measured in wafer starts per week per Fab. In addition, this increase has been achieved across more Fabs each generation. The net result is a greater than 20x

increase in normalized die output in early ramp over the past six generations. Figure 2 illustrates the rapid increase in yield-learning trends over the last seven generations. The graph shows defect learning rates (the y-axis is the logarithm of defect density, so lower is better) for Intel technologies from the start of process development through initial production. There are three key points in this data. First, the elapsed time from the start of development to the point of high yield is decreasing with subsequent technology generations. Second, the inflection point, where yield learning slows down, is occurring at higher yields with subsequent generations. Finally, the time between new process introductions is decreasing. The net result is a greater than 5x increase in normalized good die per wafer at the start of production, over the past seven generations.

These continuously increasing ramp rates and ever-improving yield-learning rates have been instrumental in maintaining Intel's leadership in the technology race, as defined by Moore's Law. There are three primary methods that enable rapid yield learning and manufacturing ramp. The first is predictive in-line metrology to shorten the cycle time for yield improvement feedback. The second is designing the process for manufacturability and performance, including using advanced process control and developing new materials. The final method is the Copy Exactly! process for transfer and ramp. The first two methods are discussed in detail elsewhere [4]. This paper focuses on Copy Exactly!.

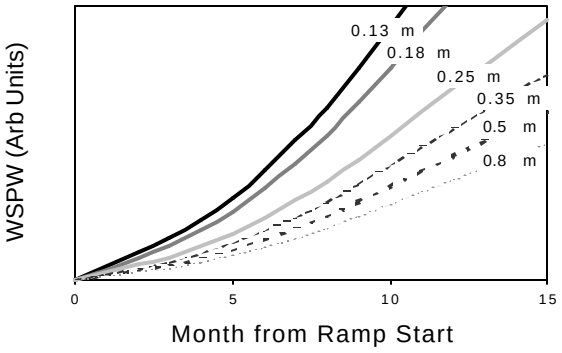

Figure 1: Intel high-volume production ramp rates

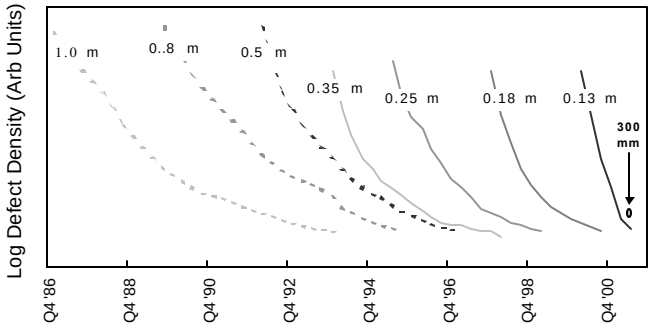

Figure 2: Intel defect density trends

COPY EXACTLY!

Up to Intel's 1 m process technology, die yields were becoming increasingly harder to match as processes were transferred from development to manufacturing facilities. During the 1 m process transfer, the first production Fab attempted to copy the development Fab closely while the second and third Fabs instituted changes (intended to be process improvements) during transfer. The results, shown in Figure 3, are striking. The so-called improvements actually resulted in an up to 10x *reduction* in die per wafer compared to the development Fab and first production Fab. This phenomenon led to the development of the Copy Exactly! methodology.

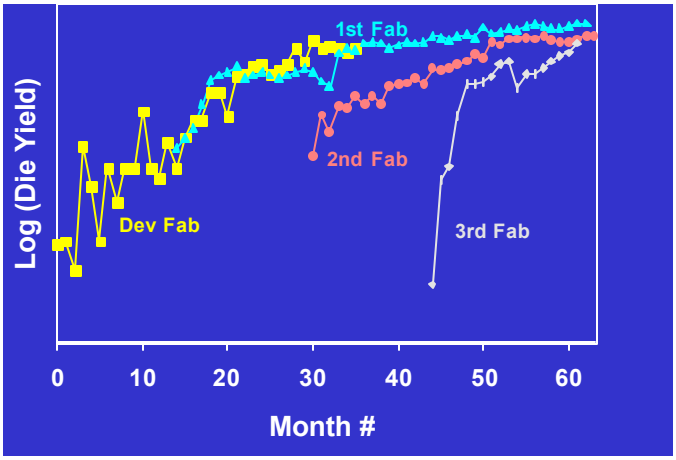

Figure 3: The birth of Copy Exactly!

The current Copy Exactly! methodology used at Intel is shown in Figure 4. The key principle behind Copy Exactly! is that Fabs running a given process technology strive to be matched in every respect except where prohibited by hard barriers. Physical inputs, such as chemical sources and purities, facilities, and hookups are all derived from the same specifications. Likewise,

equipment configurations and process recipes are matched exactly, and monitors that predict yield, reliability, and performance are all matched to within 1.5 . Once matched, changes are coordinated through cross-Fab joint engineering teams. Audits of equipment configurations and process monitors are routinely done to ensure ongoing matching. High-level tactical and strategic changes are executed in all Fabs under joint engineering and management structures.

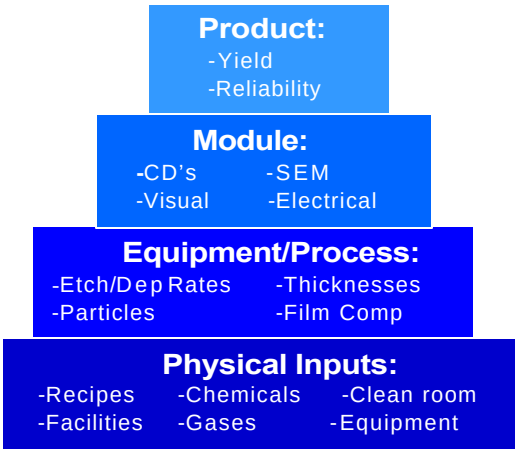

Figure 4: Current Copy Exactly! methodology

Figure 5 shows the benefit this methodology has brought since the 0.5 m technology generation. In contrast to the range of die yields observed in the 1 m generation without Copy Exactly!, every generation from the 0.5 m generation to the most recent 0.13 m generation has seen multiple Fabs started with matched die yields.

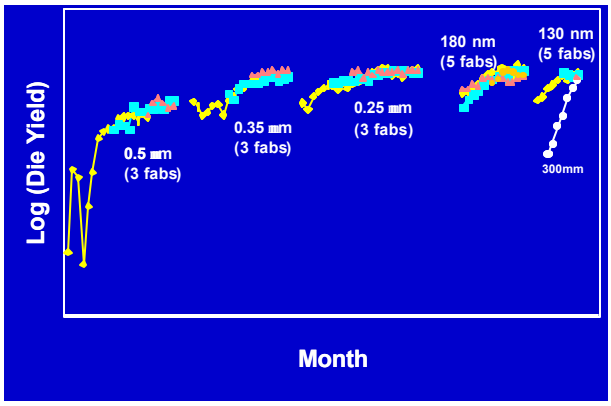

Figure 5: Die yield matching with Copy Exactly!

OVERVIEW OF INTEL'S 0.13m LOGIC TECHNOLOGY

Most recently, Intel led the industry in 2001 with the volume manufacturing ramp of a 0.13 m CMOS technology featuring 70nm dual V_t transistors, copper and low k (dielectric constant) interconnects and 2 m^2 SRAM cell sizes [1]. Table 1 summarizes the design rules for this process technology. Figures 6 and 7 illustrate Pentium III processor die size and show the relative performance between this technology and the previous 0.18 m process generation. The transition from 0.18 m to 0.13 m process technology yields a greater than 40% increase in product frequency.

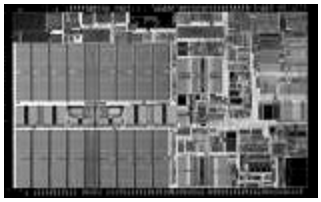

Figure 6: Pentium® III die on 0.13m process

Layer	Pitch (nm)	Thickness(nm)	Aspect Ratio
Isolation		345 450	-
Polysilicon	319	160	-
Metal 1	293	280	1.7
Metal 2, 3	425	360	1.7
Metal 4	718	570	1.6
Metal 5	1064	900	1.7
Metal 6	1143	1200	2.1

Table 1: Intel's 0.13m CMOS design rules

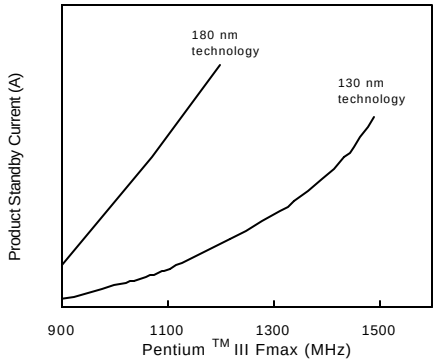

Figure 7: Pentium® III performance on 0.13µm and 0.18µm processes

300mm Wafer-Size Conversion

Intel chose the 0.13 m generation to make the wafer size change from 200mm (8") to 300mm (12"). This wafer size increase is part of an ongoing evolution beginning over 30 years ago with 1" wafers. The key driver for wafer size increase is cost reduction. The larger wafers provide a 2.25x increase in area and, due to the rectangular die size, an even larger increase in die per wafer. Manufacturing costs per wafer scale at less than this rate, so there is an overall reduction in cost per die at the larger wafer size.

The 300mm wafer size also brought a unique challenge. For the first time, the wafer size had grown large enough to pose an ergonomic hazard. A full lot of 300mm wafers weighs 18 lbs., and manual handling of 300mm wafers is prohibited due to ergonomic risks. In contrast, a full lot of 200mm wafers weighs 8 lbs. and is much smaller than a lot of 300mm wafers. 200mm wafer lots are routinely handled manually. The requirement for automated and mechanically-assisted wafer handling posed by the 300mm wafer size translates into longer cycle times for routine Fab tasks and ultimately translated into overall delays during process development.

The principal issue, however, in wafer size conversions is that the equipment set and process recipes must be completely changed to support the larger wafer. 300mm process equipment was selected using a rigorous and data-based approach. Similarity to the existing 200mm toolset was not a major factor during equipment selection: technical capability, cost, extendibility to future technologies, and productivity were. This selection process delivered a highly capable and productive toolset that could be reused for future technologies, but it drove changes away from well-characterized but less productive toolsets that had been operating, in some cases, for many years in Intel Fabs. A state-of-the-art process such as Intel's 0.13 m process has several-hundred process steps using 50-100 unique process tools. For every step, recipes must be rewritten to accommodate the larger wafers, but the higher-level goal is that the 300mm process must be essentially identical to the 200mm process in performance, reliability, and yield. With a completely new toolset and recipes that could not be simply copied or scaled, Intel faced a huge challenge in matching outputs between its 200mm and 300mm technologies. To meet the challenge, the Copy Exactly! process was adapted. This adaptation is described in the next section.

300MM COPY EXACTLY!

The development of the 300mm 0.13 m process used a modified Copy Exactly! process. Because the equipment was, by definition, different, and facility changes had to be made to accommodate the new equipment and new wafers, many of the physical inputs could not be matched. Figure 8 illustrates this. At the physical input level, recipes, facilities, equipment, and cleanroom were all not matched to 200mm. At the equipment and process level, many characteristics could not be matched because the tools either operated in different regimes from their 200mm equivalents or were based on different operating principles altogether.

However, to achieve matched output at the highest level, matching to 200mm was very extensive in other areas. To a large extent, chemicals and gases were matched, in some cases sharing a common distribution system with 200mm. Recipes were optimized for 300mm based on scaling 200mm recipes wherever possible, matching tool-level outputs to 200mm wherever possible, and always matching critical inputs to tools. "Critical inputs" are defined as those that have an impact on the wafer beyond their intended process step. For example, temperature in a thermal oxidation operation is considered a critical input because, in addition to modulating the film properties (the intended process step), temperature may also have an unintended impact on dopant diffusion and activation.

Critical outputs such as film thickness, profiles, and electrical properties were matched to 200mm within 1.5 . Variability was targeted to be equivalent or better than 200mm. The results of applying this methodology are presented in the next section.

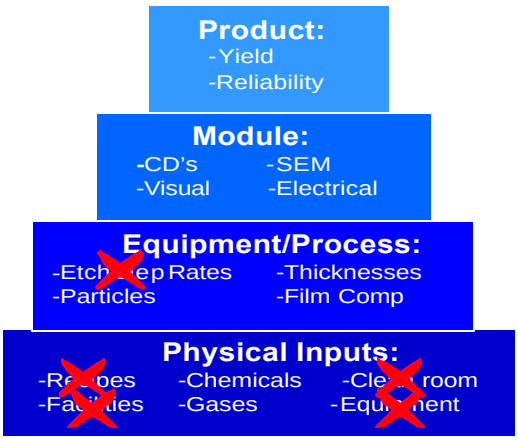

Figure 8: 200mm-to-300mm Copy Exactly! methodology

RESULTS

We now review several key metrics of our 300mm 0.13 m process and compare them to the 200mm process. We begin with module-level data, characterizing the matching of specific tools or subsets of the overall process. We then report matching data on transistor and Pentium 4 processor product performance, yield, and reliability. The data shown are a representative sample of all such indicators. In general, all data are matched between 200mm and 300mm to a similar degree. Across the board, the data show excellent matching between the 200mm and 300mm 0.13 m processes.

Module-Level Matching

Figure 9 shows within-wafer matching for a representative in-line key monitor. Shown here are 200mm and 300mm wafer maps of gate-oxide thickness. The data show that 300mm wafers have slightly better within-wafer gate-oxide thickness variation than 200mm wafers.

Figure 10 shows cumulative distributions for back-end Via resistances for 200mm and 300mm wafers. Via resistance is an integrated measure of interconnect electrical performance. As the data show, 200mm and 300mm Via resistances are closely matched.

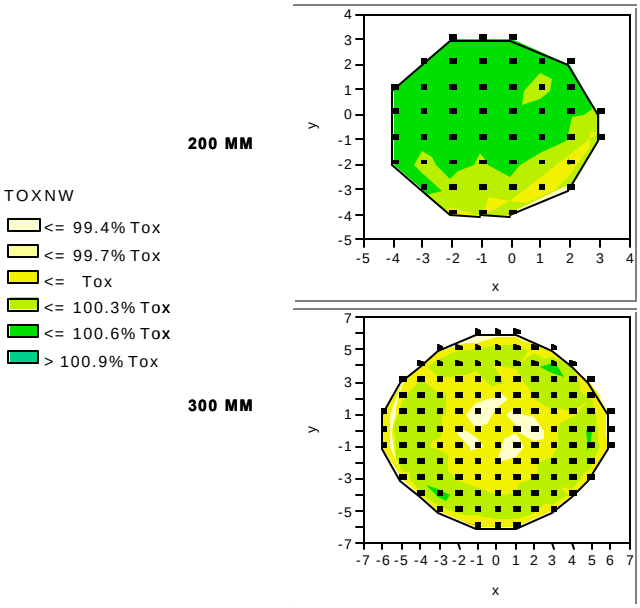

Figure 9: 200mm/300mm within-wafer gate oxide

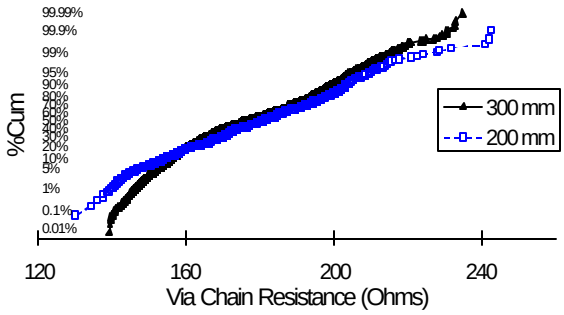

Figure 10: 200mm/300mm Via resistance distribution

Figures 11(a) and 11(b) show Transmission Electron Microscope (TEM) cross-sections of 200mm and 300mm gate electrodes. These are approximately identical, non-minimum gate-length transistors. Profiles and critical film thicknesses are well matched. Slight differences in the film conformality and interfaces are evident. These are unavoidable differences caused by configuration differences between the 200mm and 300mm tools.

Figure 12 shows a TEM cross-section of the complete 6-layer interconnect system. Profiles and thickness are virtually identical between 200mm and 300mm.

Pentium is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

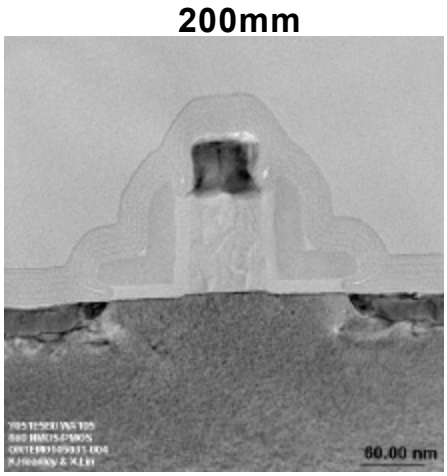

Figure 11(a): 200mm gate electrode TEMs

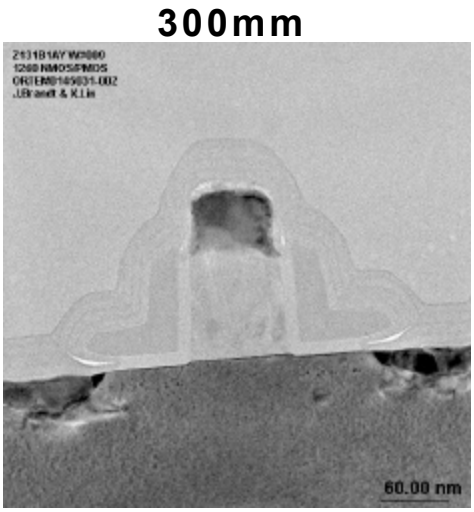

Figure 11(b) : 300mm gate electrode TEMs

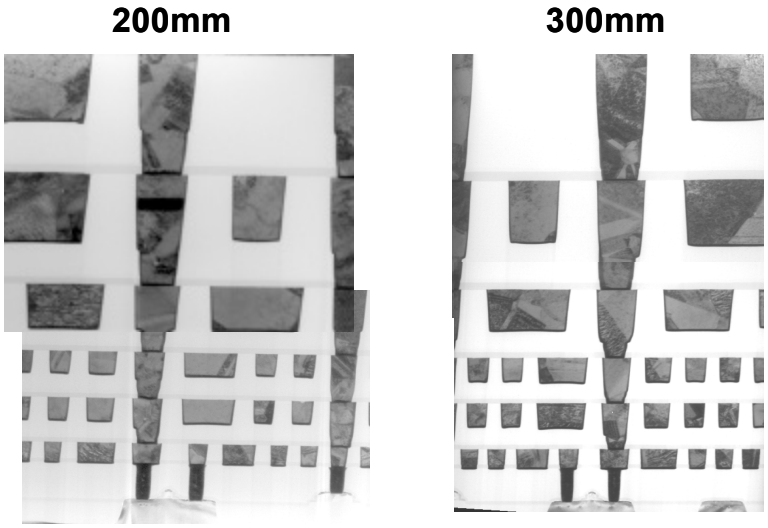

Figure 12: 200mm/300mm interconnect TEMs

Performance Matching

Figure 13 shows a basic transistor matching graph between 200mm and 300mm. Saturated drive current (Idsat) is plotted against off-state leakage (Ioff) for both 200mm and 300mm NMOS and PMOS transistors. The data show that the 200mm and 300mm devices are perfectly matched across a wide range of Ioff.

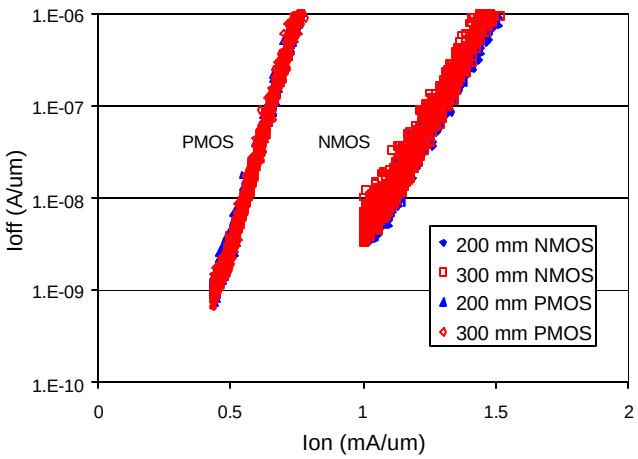

Figure 13: 200mm/300mm transistor Ion/Ioff

Figure 14 shows a circuit-level matching metric. The graph is a cumulative distribution of ring oscillator test circuit frequencies on 200mm and 300mm wafers. Again, the data indicate that the circuit operating frequencies are perfectly matched.

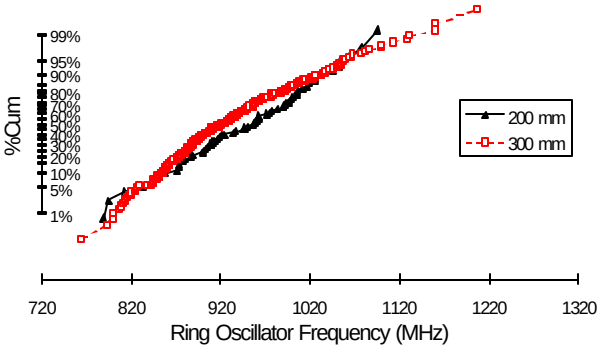

Figure 14: 200mm/300mm ring oscillator circuit frequency

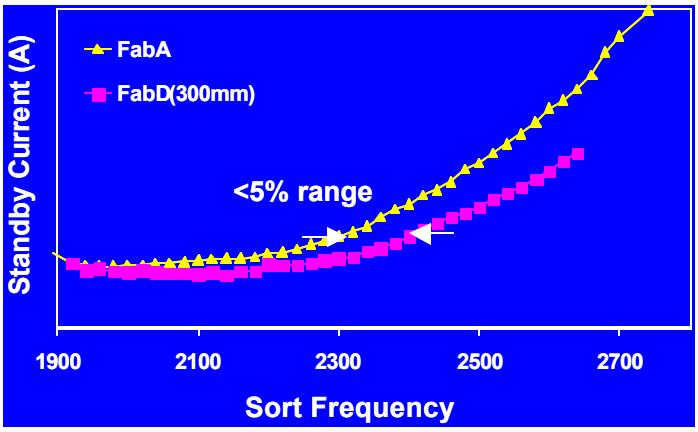

Figure 15: Pentium® 4 sort frequency vs. standby current

Finally, Figure 15 shows a normalized performance comparison for the Pentium® 4 product. The graph shows sort frequency graphed against product standby current. The 300mm product speed is within 5% of the reference 200mm population, matched to within normal variability.

Yield Matching

Figure 16 shows normalized die yield for 300mm and 200mm as a function of time. 300mm die yield at the start of development is lower than 200mm, which is shown starting after initial ramp. Rapid yield learning, facilitated by the ability to copy 200mm learning, enabled steadily improving die yields to the point where 200mm and 300mm die yields are matched at the point of the 300mm initial ramp.

Reliability Matching

Figure 17 shows a key transistor reliability metric, gate-oxide time-to-breakdown. The data are shown as a normalized distribution function of time-dependent dielectric breakdown (TDDB) in seconds. Both 200mm and 300mm are well matched in gate-oxide reliability.

Figure 18 shows a key interconnect reliability metric, electromigration fail rate. The data are shown as a normalized distribution function of time-to-fail. Again, both 200mm and 300mm are well matched.

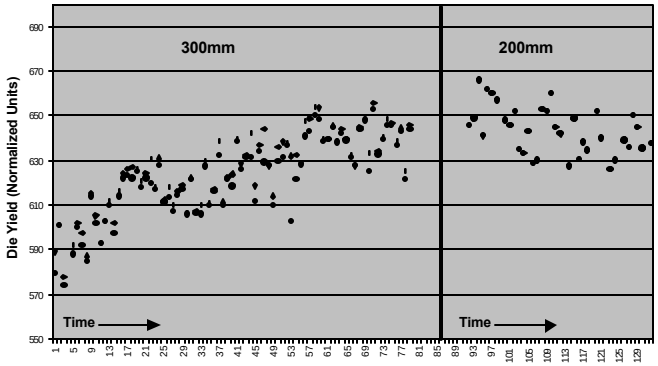

Figure 16 : 200mm/300mm normalize die yield

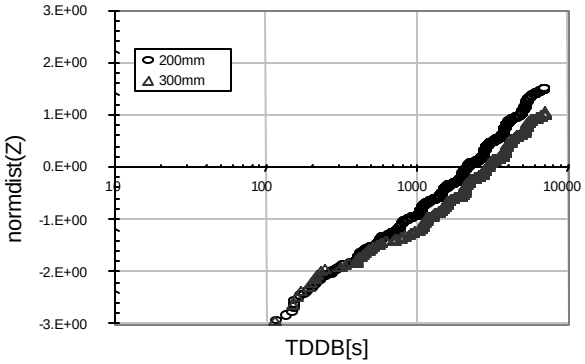

Figure 17: Gate oxide 200mm/300mm time-to-fail

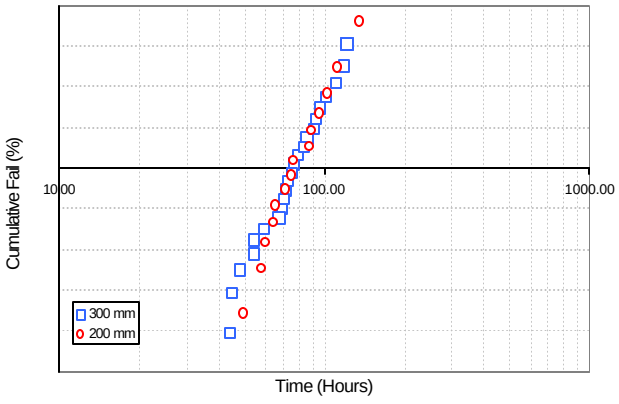

Figure 18: Electromigration cumulative fail rate

CONCLUSION

For over 35 years, Moore's Law has set a rapid pace for progress in the semiconductor industry. With the continuously increasing technical challenges for silicon technology development, increasingly rapid yield learning and volume manufacturing ramp rate have been instrumental in maintaining Intel's technology leadership.

In this paper, we discussed the implementation of Intel's industry-leading 0.13 m logic technology on the 300mm wafer size and associated process equipment. The 0.13 m process has been ramped to volume production in multiple factories and on both 200mm and 300mm production lines at record yields, quality, and ramp rate. Rapid development of Intel's first 300mm wafer-size technology, well matched to the 200mm state-of-the-art process, is a critical milestone for future competitiveness. The adaptation of proven Copy Exactly! methods is the key element that enabled successful conversion to the 300mm wafer size and sets the stage for Intel's continued leadership in the semiconductor industry.

ACKNOWLEDGMENTS

The authors thank the many outstanding engineers and technicians in the Logic Technology Development and 300mm OHT organizations who are responsible for developing, ramping, and transferring Intel's P1260 logic technology.

REFERENCES

[1] Tyagi, S., et. al., "A 130nm Generation Logic Technology Featuring 70nm Transistors, Dual Vt Transistors and 6 layers of Cu Interconnects," *IEDM Technical Digest*, December, 2000, pp. 567-570.

[2] Moore, G.E. "Cramming More Components Onto Integrated Circuits," *Electronics Magazine*, Vol. 8, April, 1965, pp. 114-117.

[3] Schaller, R.R. "Moore's Law: Past, Present, and Future," *IEEE Spectrum*, Vol. 34, Issue 6, pp. 52-59.

[4] Gasser, Jr., R.A., "Yield Learning and Volume Manufacturing of High-Performance Logic Technologies on 200mm and 300mm Wafers," *IEDM Technical Digest*, December, 2001, pp. 599-601.

AUTHORS' BIOGRAPHIES

Sanjay Natarajan is a process integration group leader in Intel's Logic Technology Development organization. He joined Intel in 1993 and has held numerous positions in both factory automation and process integration. He is presently responsible for transistor integration for Intel's 65nm CMOS process technology. Prior to this, he led transistor integration for Intel's first 300mm process technology, a 0.13 m CMOS process matched to Intel's 200mm technology. He received his B.S., M.S., and Ph.D. degrees from Carnegie Mellon University, all in Electrical Engineering. His e-mail is Sanjay.Natarajan@intel.com.

Melton Bost is a process integration group leader in Intel's Logic Technology Development organization. He joined Intel in 1987 and has worked primarily in back- end process integration since that time. He is presently responsible for the P1260 backend process technology. He received his B.S.E. degree from Duke University in 1978, his M.S.E. degree from Stanford University in 1979, both in Materials Science, and his Ph.D. degree from Colorado State University in 1987 in Electrical Engineering. He is the author of numerous technical papers and holds six patents. His e-mail is Melton.Bost@intel.com.

Derek Fisher is the P1260 Yield Group Leader in Intel's Logic Technology Development organization. Since 1994, he has been responsible for Defect Metrology Roadmaps and participated in Fab startup and process transfer activities from P854 to P1260. He joined Intel in Ireland in 1991 and worked on P652 to P852 transfer and yield improvement. Prior to that, he worked in manufacturing and process development roles for National Semiconductor, Philips Research, and Motorola. He graduated from Strathclyde University in 1983. His e-mail is Derek.G.Fisher@intel.com.

David Krick is the equipment startup coordinator for Intel's Logic Technology Development organization. He joined Intel in 1989 as a process engineer and has held numerous process engineering and management positions. He is currently responsible for coordinating the startup of Intel's newest 300mm technology development facility, D1D. Prior to this, he managed the successful startup of D1C, Intel's first 300mm factory. He holds three patents. Krick received his B.S. degree in Electrical Engineering and his M.S. degree in Engineering Science, both from Pennsylvania State University. His e-mail is David.T.Krick@intel.com.

Chris Kenyon is a lithography group leader in Intel's Logic Technology Development organization. He joined Intel in 1996 and has worked primarily on Intel's gate patterning process since that time. He is currently responsible for developing the gate patterning process for the 90nm CMOS node. He received his B.A. degree from Princeton University in 1990 and his Ph.D. degree from Caltech in 1996 in Physical Chemistry. His e-mail is Chris.Kenyon@intel.com.

Chris Kardas is a process engineering group leader in Intel's Logic Technology Development organization. Since joining Intel in 1984, he has worked on a variety of Etch modules where he holds a patent. Most recently he has been working in the area of lithography and is presently responsible for P1260 and P1212 front-end layers. He received his B.S.E.E. degree from the University of Illinois in 1984. His e-mail is Chris.Kardas@Intel.com.

Chris Parker is a senior process engineer in Intel's Logic Technology Development organization. He joined Intel in 1998 and has worked on front-end oxidation and gate development. He is presently responsible for the development of alternative gate dielectric processes for 300mm P126x technologies. He received his B.S.E.E. degree from Auburn University and M.S. and Ph.D. degrees in Electrical Engineering from North Carolina State University. His e-mail is Chris.PTD.Parker@intel.com.

Robert A. Gasser, Jr. is Vice President, Technology & Manufacturing Group, and Director, Components Research. He is responsible for research and development of process technologies used to build future Intel logic devices that will be in production five to ten years from now. He joined Intel in 1982 as a Technology Evaluation Engineer. Most recently, Gasser was responsible for the development of Intel's 0.18-micron logic process technology (P858) and 300mm, 130nm logic technology (P1260) and 90nm logic technology (P1262). Gasser received his B.A. degree in Physics from Reed College in 1980. He received his M.S. degree in Materials Science from Stanford University in 1982. Gasser has written numerous technology papers and holds five patents. His e-mail is Bob.Gasser@intel.com.

For further information visit:

developer.intel.com/technology/itj/index.htm